arXiv:1610.09908v1 [cs.CV] 31 Oct 2016

Joint Large-Scale Motion Estimation and Image Reconstruction

Hendrik Dirks*

November 1, 2016

Abstract

This article describes the implementation of the joint motion estimation and image reconstruction framework presented by Burger, Dirks and Schönlieb and extends this framework to large-scale motion between consecutive image frames.
The variational framework uses displacements between consecutive frames based on the optical flow approach to improve the image reconstruction quality on the one hand and the motion estimation quality on the other. The energy functional consists of a data-fidelity term with a general operator that connects the input sequence to the solution, it has a total variation term for the image sequence and is connected to the underlying flow using an optical flow term. Additional spatial regularity for the flow is modeled by a total variation regularizer for both components of the flow.
The numerical minimization is performed in an alternating manner using primal-dual techniques. The resulting schemes are presented as pseudo-code together with a short numerical evaluation.

1 Introduction

Image reconstruction and motion estimation are important tasks in image processing. Such problems arise for example in modern medicine, biology,

*Institute for Computational and Applied Mathematics and Cells in Motion Cluster of Excellence, University of Münster, Orléans-Ring 10, 48149 Münster, Germany, Email: hendrik.dirks@wwu.de

chemistry or physics, where even the smallest objects are observed by high resolution microscopes. To characterize the dynamics involved in such data, velocity fields between consecutive image frames are calculated. This is challenging, since the recorded images often suffer from low resolution, low contrast, different gray levels and noise. Methods that simultaneously denoise the recorded image sequence and calculate the underlying velocity field offer new opportunities, since both tasks may endorse each other.
The ansatz from [6] aims at reconstructing a given sequence u of images and calculating flow fields $\boldsymbol{v}$ between subsequent images at the same time. For given measurements $f = Ku$ this can be achieved by minimizing the variational model

$$\int_0^T \frac{1}{2} \|K_t u(\cdot,t) - f(\cdot,t)\|_2^2 + \alpha \mathcal{R}(u(\cdot,t)) + \beta \mathcal{S}(\boldsymbol{v}(\cdot,t)) + \|\rho(u,\boldsymbol{v})\|_1 dt \quad (1)$$

with respect to u and $\boldsymbol{v}$ simultaneously. The denoising part is based on the ROF model [17]. The first part $\|Ku - f\|_2^2$ connects the input data f with the image sequence u via a linear operator K. Depending on the application, K may model the cutting out of a subset $\Sigma \subset \Omega$ for inpainting, a subsampling for super resolution, a blur for deconvolution or a Radon transform for computed tomography. Additional a-priori information about the structure of u respectively $\boldsymbol{v}$ can be incorporated into each frame via the regularization terms $\mathcal{R}(u(\cdot,t))$ and $\mathcal{S}(\boldsymbol{v}(\cdot,t))$, while their significance is weighted using α and β. Finally, flow field and images are coupled by a model for the underlying dynamics, e.g. the optical flow approach $\rho(u,\boldsymbol{v}) = u_t + \nabla u \cdot \boldsymbol{v}$ (see Section 1.2).

1.1 Notations

Let us for the following assume a discrete rectangular domain $\Omega \subset \mathbb{R}^2$ and sequences of recorded (noisy) images $f^1, \dots, f^n : \Omega \to \mathbb{R}$, clean images $u^1, \dots, u^n : \Omega \to \mathbb{R}$ and motion fields between subsequent frames $\boldsymbol{v}^1, \dots, \boldsymbol{v}^{n-1} : \Omega \to \mathbb{R}^2$.

1.2 Optical Flow

In the optical flow problem we want to estimate the flow field $\boldsymbol{v}^i = (v^{i,1}, v^{i,2})^T$ that describes the displacement between each of the subsequent images u^i and

u^{i+1}. For this sake it is common to use the brightness constancy assumption

$$u^{i+1}(x + \boldsymbol{v}^i(x)) - u^i(x) = 0, \quad i = 1, \dots, n-1, x \in \Omega \tag{2}$$

to derive a connection between image intensities and the underlying flow. Due to the non-linearity (in terms of $\boldsymbol{v}^i(x)$) of this formulation, one often linearizes the first term and arrives at

$$\boldsymbol{v}^i \cdot \nabla u^{i+1}(x) + u^{i+1}(x) - u^i(x) = 0. \tag{3}$$

Unfortunately, due to the Taylor expansion, equation (3) is only valid for small displacements. Another way, which in theory may handle displacements of arbitrary magnitude, is to use a given flow field $\tilde{\boldsymbol{v}}^i$ and use a Taylor expansion of $u^{i+1}(x + \boldsymbol{v}^i(x))$ with respect to a known flow field $\tilde{\boldsymbol{v}}^i$. The result is again a linear equation for $\boldsymbol{v}^i$ which requires evaluations of the input images at a shifted domain $x + \tilde{\boldsymbol{v}}^{\boldsymbol{i}}$

$$(\boldsymbol{v}^i - \tilde{\boldsymbol{v}}^i) \cdot \nabla u^{i+1}(x + \tilde{\boldsymbol{v}}^i) + u^{i+1}(x + \tilde{\boldsymbol{v}}^i) - u^i(x) = 0. \tag{4}$$

By defining $\tilde{u}^{i+1} := u^{i+1}(x + \tilde{\boldsymbol{v}}^i)$ the left side of equation (4) becomes

$$\rho(\boldsymbol{v}^i, u^i, u^{i+1}) := (\boldsymbol{v}^i - \tilde{\boldsymbol{v}}^i) \cdot \nabla \tilde{u}^{i+1} + \tilde{u}^{i+1} - u^i. \tag{5}$$

Both, formulation (3) and (4), state only one equation per point for the two unknown components of $\boldsymbol{v}$ and, consequently, the problem is underdetermined. To overcome this, the optical flow formulation can be used as a data fidelity in a variational model together with an isotropic total variation term on each of the two flow components to ensure spatial regularity. This leads to the following variational problem for calculating the optical flow of a video sequence $u^1, \dots, u^n$:

$$\operatorname*{arg\,min}_{\boldsymbol{v}=\boldsymbol{v}^1,\dots,\boldsymbol{v}^{n-1}} \sum_{i=1}^{n-1} \|\rho(\boldsymbol{v}^i, u^i, u^{i+1})\|_1 + \alpha \sum_{j=1}^{2} \left\|\nabla v^{i,j}\right\|_{1,2}. \tag{6}$$

This variational approach to optical flow was first introduced by Horn and Schunck [10] and continuously improved [1, 14]. Further details about the so-called L^1-TV optical flow model above can be found e.g. in [23, 24].

1.3 Image Reconstruction

Nowadays, variational models for image reconstruction have become very popular. One of the most famous models, introduced by Rudin, Osher and Fatemi in 1992 [17], is the total variation (TV) model, where the authors couple a L^2 data fidelity term with a total variation regularization. Data-term and regularizer in the ROF model match with the first two terms model (1). The TV-regularization results in a denoised image with cartoon-like features. This model has also been adapted to image deblurring [22], inpainting [19], super resolution [12, 21] and tomographic reconstruction [18, 11]. We collectively call these image reconstruction models.
From our point of view, the problem of motion estimation is directly connected to the underlying image sequence and, hence, requires accurate input images u^i. Unfortunately, for many practical applications, only noisy variants f^i of u^i can be recorded. Similar to the optical flow problem, a variational formulation can be used to reconstruct u^i from f^i. Let us assume that f^i is a degraded version of u^i corrupted by Gaussian noise. Then, the ROF model

$$\underset{u=u^1,\dots,u^n}{\arg\min} \sum_{i=1}^{n} \frac{1}{2}\|A^i u^i - f^i\|_2^2 + \alpha\|\nabla u^i\|_{1,2} \tag{7}$$

can be used to reconstruct each u^i from f^i. Here, A^i represents some linear operator and could be the identity (denoising), subsampling (zooming) or the Radon transform (computed tomography).

1.4 Joint Model

As already mentioned, motion estimation should be done on noise-free images, so generally one first denoises the image sequence using (7) and afterwards estimates the underlying velocity fields using (6). In [6] it has been shown that a joint model that simultaneously recovers an image sequence and estimates motion offers a significant advantage towards subsequently applying both methods. The following joint model

$$\underset{u,\boldsymbol{v}}{\arg\min} \int_0^T \frac{1}{2}\,\|Au - f\|_2^2 + \alpha\,\|\nabla u\|_{1,2} + \gamma\,\|u_t + \nabla u \cdot \boldsymbol{v}\|_1 + \beta \sum_{j=1}^{2} \|\nabla v_j\|_{1,2}\,dt, \tag{8}$$

respectively its time-discrete counterpart

$$\underset{\substack{u=u^1,\dots,u^n\\ \boldsymbol{v}=\boldsymbol{v}^1,\dots\boldsymbol{v}^{n-1}}}{\arg\min} \sum_{i=1}^{n} \frac{1}{2} \left\|A^i u^i - f^i\right\|_2^2 + \alpha \left\|\nabla u^i\right\|_{1,2} + \sum_{i=1}^{n-1} \gamma \left\|u_t^i + \nabla u^i \cdot \boldsymbol{v}^i\right\|_1 + \beta \sum_{j=1}^{2} \left\|\nabla v^{i,j}\right\|_{1,2}, \tag{9}$$

was proposed. For both, image sequence and velocity field, the respective total variation is used as a regularizer and the classical optical flow formulation from (3) connects image sequence and velocity field. From the perspective of image reconstruction the optical flow constraint acts as an additional temporal regularizer along the calculated motion fields $\boldsymbol{v}$. In [4] the existence of a minimizer for model (8) under certain regularity assumptions for $\boldsymbol{v}$ and $\nabla \boldsymbol{v}$ has been shown. Despite the existence of a minimizer, calculating it is numerically challenging. Problems arise from the non-convexity and non-linearity of the optical flow term, from the non-differentiability of the L^1–norm and finally due to several linear operators acting on u and $\boldsymbol{v}$.
The main drawback in terms of practical applications for (9) is the restriction to displacements of small magnitude, which cannot be expected in practical applications. By exchanging the classical optical flow constraint with the beneficial linearization from (4) we create a model capable of handling large-scale displacements. The joint large-scale motion estimation and image reconstruction model then reads

$$\underset{\substack{u=u^1,\dots,u^n\\ \boldsymbol{v}=\boldsymbol{v}^1,\dots\boldsymbol{v}^{n-1}}}{\arg\min} \sum_{i=1}^{n} \frac{1}{2} \|A^i u^i - f^i\|_2^2 + \alpha \left\|\nabla u^i\right\|_{1,2} + \sum_{i=1}^{n-1} \|\rho(\boldsymbol{v}^i, u^i, u^{i+1})\|_1 + \beta \sum_{j=1}^{2} \left\|\nabla v^{i,j}\right\|_{1,2}. \tag{10}$$

By setting $\tilde{\boldsymbol{v}}^i = 0$ in the equation $\rho(\boldsymbol{v}^i, u^i, u^{i+1})$ we obtain the original energy.

2 Numerical Scheme and Implementation

We propose, referring to [6], a minimization scheme which alternatingly fixes u and $\boldsymbol{v}$ and minimizes the energy for the other variable. The corresponding problems for static $u_k^1, \dots, u_k^n$ in equation (11) resp. $\boldsymbol{v}_{k+1}^1, \dots \boldsymbol{v}_{k+1}^{n-1}$ in equation

(12) read

$$\boldsymbol{v}_{k+1} = \underset{\boldsymbol{v}=\boldsymbol{v}^1,\dots\boldsymbol{v}^{n-1}}{\arg\min} \sum_{i=1}^{n-1} \|\rho(\boldsymbol{v}^i, u_k^i, u_k^{i+1})\|_1 + \beta \sum_{j=1}^{2} \left\|\nabla v^{i,j}\right\|_{1,2}, \tag{11}$$

$$u_{k+1} = \underset{u=u^1,\dots,u^n}{\arg\min} \sum_{i=1}^{n} \frac{1}{2}\|A^i u^i - f^i\|_2^2 + \alpha\|\nabla u^i\|_1 + \sum_{i=1}^{n-1} \|\rho(\boldsymbol{v}_{k+1}^i, u^i, u^{i+1})\|_1, \tag{12}$$

where the letter k shall denote the iteration number. To start with the minimization problem in (11), let us denote that in this formulation the evaluation of u_k^{i+1} on intermediate grid points $x + \tilde{\boldsymbol{v}}$ is required. These evaluations can be generated by an interpolation scheme (linear, cubic etc.). Moreover, the linearization of the brightness constancy assumption requires a flow field $\tilde{\boldsymbol{v}}$ close to $\boldsymbol{v}$. This can be generated by an iterative coarse-to-fine approach, which solves the problem on subsampled versions of u^i first and using the upscaled result as initial value for the next finer version. Moreover, on each level several so-called warpings are performed, where the problem is solved for some initial $\tilde{\boldsymbol{v}}$ and the solution is used in the next step (see Section 2.1.1 and [23, 24] for details).
Treating the minimization problem in equation (12) next, this problem incorporates the term

$$\rho(\boldsymbol{v}_{k+1}^i, u^i, u^{i+1}) = (\boldsymbol{v}_{k+1}^i - \tilde{\boldsymbol{v}}_{k+1}^i) \cdot \nabla\tilde{u}^{i+1} + \tilde{u}^{i+1} - u^i \tag{13}$$

Let us recall that in the minimization process for $\boldsymbol{v}^i$, a warping scheme is applied, which creates subsequent versions of $\boldsymbol{v}$ and $\tilde{\boldsymbol{v}}$ which can be assumed to converge to a static quantity. Consequently, the difference $\boldsymbol{v}_{k+1}^i - \tilde{\boldsymbol{v}}_{k+1}^i$ becomes arbitrarily small, so it can be neglected in (13) and the term simplifies to $\|\tilde{u}^{i+1} - u^i(x)\|_1$. We want to underline that for evaluating $\tilde{u}^{i+1}$ the same interpolation scheme as in the $\boldsymbol{v}$-problem has to be used to ensure numerical consistency.

2.1 Problem in $\boldsymbol{v}$

The optical flow problem has no time-correspondence in $\boldsymbol{v}$ and thus reduces to $(n-1)$ subproblems of the form

$$\underset{\boldsymbol{v}}{\arg\min} \|\boldsymbol{v} \cdot \nabla\tilde{u} + \tilde{u}_t\|_1 + \beta \sum_{j=1}^{2} \left\|\nabla v^j\right\|_{1,2}, \tag{14}$$

where $\tilde{u}_t := -\tilde{\boldsymbol{v}}^i \cdot \nabla \tilde{u}^{i+1} + \tilde{u}^{i+1} - u^i$ and $\tilde{u} := \tilde{u}^{i+1}$. Applying a primal-dual algorithm [16, 5] yields the following iterative scheme:

$$
\begin{aligned}
\boldsymbol{y}_1^{k+1} &= \frac{\boldsymbol{y}_1^k + \sigma_1^v \nabla \bar{v}_1^k}{\max(1, \frac{\|\boldsymbol{y}_1^k + \sigma_1^v \nabla \bar{v}_1^k\|_2}{\alpha})}, \\
\boldsymbol{y}_2^{k+1} &= \frac{\boldsymbol{y}_2^k + \sigma_2^v \nabla \bar{v}_2^k}{\max(1, \frac{\|\boldsymbol{y}_2^k + \sigma_2^v \nabla \bar{v}_2^k\|_2}{\alpha})}, \\
y_3^{k+1} &= \max(-1, \min(1, y_3^k + \sigma_3^v (\nabla \tilde{u} \cdot \bar{\boldsymbol{v}}^k + \tilde{u}_t))), \\
v_1^{k+1} &= v_1^k - \tau_1^v (\nabla^T \cdot \boldsymbol{y}_1^{k+1} + \tilde{u}_x y_3^{k+1}), \\
v_2^{k+1} &= v_2^k - \tau_2^v (\nabla^T \cdot \boldsymbol{y}_2^{k+1} + \tilde{u}_y y_3^{k+1}), \\
\bar{\boldsymbol{v}}^{k+1} &= 2\boldsymbol{v}^{k+1} - \boldsymbol{v}^k,
\end{aligned}
$$

where vector-valued dual variables $\boldsymbol{y}_1, \boldsymbol{y}_2$, a regular dual variable y_3 and step sizes $\sigma_1^v, \sigma_2^v, \sigma_3^v, \tau_1^v$ and τ_2^v were introduced. As a convergence criterion we use the primal-dual residual r_v^{pd} from [9]. Applied to this problem, the residual can be calculated as

$$
\begin{aligned}
p_1 &= \|\frac{v_1^k - v_1^{k+1}}{\tau_1^v} - \nabla^T \cdot (\boldsymbol{y}_1^{k+1} - \boldsymbol{y}_1^k) - \tilde{u}_x (y_3^{k+1} - y_3^k)\|_1, \\
p_2 &= \|\frac{v_2^k - v_2^{k+1}}{\tau_2^v} - \nabla^T \cdot (\boldsymbol{y}_2^{k+1} - \boldsymbol{y}_2^k) - \tilde{u}_y (y_3^{k+1} - y_3^k)\|_1, \\
d_1 &= \|\frac{\boldsymbol{y}_1^k - \boldsymbol{y}_1^{k+1}|}{\sigma_1} - \nabla (v_1^k - v_1^{k+1})\|_1, \\
d_2 &= \|\frac{\boldsymbol{y}_2^k - \boldsymbol{y}_2^{k+1}|}{\sigma_2} - \nabla (v_2^k - v_2^{k+1})\|_1, \\
d_3 &= \|\frac{y_3^k - y_3^{k+1}|}{\sigma_3} - \nabla \tilde{u} \cdot (\boldsymbol{v}^k - \boldsymbol{v}^{k+1})\|_1, \\
r_v^{pd} &= p_1 + p_2 + d_1 + d_2 + d_3.
\end{aligned}
$$

2.1.1 Coarse-to-fine pyramid

The optical flow calculation is incorporated into a coarse-to-fine pyramid with intermediate warping steps. This is done by creating a set with $nScales$ versions of the input images with a downsampling factor $\eta \in (0, 1)$. To

create these images, the input image is convolved with a Gaussian kernel with standard variation σ_d and then extrapolated using a bicubic interpolation scheme.
The optical flow is calculated on the coarsest scale with initial $\tilde{\boldsymbol{v}} = 0$ first. The result $\boldsymbol{v}$ is filtered using a two-dimensional median-filter of size $sizeMed$. Afterwards, a bicubic interpolation scheme is taken to calculate both $\tilde{\boldsymbol{v}}$ and $\nabla\tilde{\boldsymbol{v}}$. Then we solve the problem again on the same scale. This procedure is repeated $nWarps$ times, until the result $\boldsymbol{v}$ is upscaled to the next finer level using bicubic interpolation. To improve the convergence speed, it is beneficial to also upscale the dual variables $\boldsymbol{y}_1, \boldsymbol{y}_2$ and y_3.

2.2 Problem u

The problem in u does not simplify to a series of time-independent problems, since individual frames are correlated by the flow. Consequently, the problem has to be solved in the whole space/time domain. First, we want to deduce that (12) can be rewritten in the form

$$\arg\min_u \frac{1}{2}\|\mathcal{A}u - f\|_2^2 + \alpha\|\bar{\nabla}u\|_{1,2} + \gamma\|\mathcal{W}u\|_1, \tag{15}$$

where the identities $u = (u^1, \ldots, u^n), f = (f^1, \ldots, f^n), \mathcal{A} = diag(A^1, \ldots, A^n), \bar{\nabla} = diag(\nabla, \ldots, \nabla)$ hold obviously. For the last term, note that the optical flow term from Equation (12) consists of $(n-1)$ parts of the form $\tilde{u}^{i+1} - u^i$, where $\tilde{u}^{i+1}$ is a shifted version of u^{i+1} by a velocity field $\boldsymbol{v}^i$ that can be evaluated using an interpolation scheme. This quantity can be approximated using a so-called interpolation or warping operator W^{i+1} in which each row has weights at the columns according to the interpolation method (linear, cubic, spline etc.). Thus, we have $\tilde{u}^{i+1} \approx W^{i+1}u^{i+1}$ and the last part in (12) reduces to $\|\mathcal{W}u\|_1$, where

$$\mathcal{W} = \begin{bmatrix} -I & W^2 & 0 & \ldots & \ldots & 0 \\ 0 & -I & W^3 & 0 & \ldots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & 0 & -I & W^{n-2} & 0 \\ 0 & \ldots & \ldots & 0 & -I & W^{n-1} \end{bmatrix}.$$

Applying the primal-dual algorithm [16, 5] to problem (15) we obtain the iterative scheme

$$
\begin{aligned}
y_1^{k+1} &= \frac{1}{\sigma_1^u + 1}(y_1^k + \sigma_1^u(A\bar{u}^k - f)),\\
\boldsymbol{y}_2^{k+1} &= \frac{\boldsymbol{y}_2^k + \sigma_2^u \bar{\nabla}\bar{u}^k}{\max(1, \frac{\|\boldsymbol{y}_2^k + \sigma_2^u \bar{\nabla}\bar{u}^k\|_2}{\alpha})},\\
y_3^{k+1} &= \max(-\gamma, \min(\gamma, y_3^k + \sigma_3^u \mathcal{W}\bar{u}^k)),\\
u^{k+1} &= u^k - \tau^u(A^T y_1^{k+1} + \bar{\nabla}^T \cdot \boldsymbol{y}_2^{k+1} + \mathcal{W}^T y_3^{k+1}),\\
\bar{u}^{k+1} &= 2u^{k+1} - u^k,
\end{aligned}
$$

where, similar to the problem for $\boldsymbol{v}$, regular dual variables y_1, y_3, a vector-valued dual variable $\boldsymbol{y}_2$ and step-sizes $\sigma_1^u, \sigma_2^u, \sigma_3^u$ and τ^u were introduced. Once again, we use the primal-dual residual [9] as a convergence criterion that can be calculated as follows:

$$
\begin{aligned}
p &= \|\frac{u^k - u^{k+1}}{\tau} - \mathcal{A}^T(y_1^{k+1} - y_1^k) - \bar{\nabla}^T \cdot (\boldsymbol{y}_2^{k+1} - \boldsymbol{y}_2^k) - \mathcal{W}^T(y_3^{k+1} - y_3^k)\|_1\\
d_1 &= \|\frac{y_1^k - y_1^{k+1}|}{\sigma} - A(u^k - u^{k+1})\|_1\\
d_2 &= \|\frac{\boldsymbol{y}_2^k - \boldsymbol{y}_2^{k+1}|}{\sigma} - \bar{\nabla}(u^k - u^{k+1})\|_1\\
d_3 &= \|\frac{y_3^k - y_3^{k+1}|}{\sigma} - \mathcal{W}(u^k - u^{k+1})\|_1\\
r_u^{pd} &= p + d_1 + d_2 + d_3
\end{aligned}
$$

2.3 Initial Data

Due to the fact that the initial formulation (10) of the problem is highly non-linear, our alternating minimization scheme only solves the linear sub-problems and will most likely end up in a local minimum. In our numerical evaluation, we found that initializing the image series u with reasonable values we usually end up in a lower energy. This can for example be done by solving an ROF-problem for each image u_i, i.e. setting the weight γ in (15) to zero.

Another possibility to initialize the image sequence is to neglect the flow

fields $\boldsymbol{v}$ in the whole problem. One could for example expect a smooth time-derivative u_t and solve to following modified problem

$$\arg\min_u \frac{1}{2}\|\mathcal{A}u - f\|_2^2 + \alpha\|\bar{\nabla}u\|_{1,2} + \frac{\epsilon}{2}\|u_t\|_2^2, , \tag{16}$$

where ϵ is a small weight However, depending on the application, for large displacements this may introduce smearing effects and lead to worse results than the first method.

2.4 Operator Discretization

We assume the underlying space-time grid to consist of the following set of discrete points:

$$\{(i, j, t) : i = 0, \ldots, n_x, j = 0, \ldots, n_y, t = 0, \ldots, n_t\}$$

For the sake of simplicity we write down the discrete operators applied to an example variable u respectively its adjoint $\boldsymbol{y}$. The required generalization to the appearing primal and dual variables is left to the reader. The discrete gradient is calculated using forward differences and Neumann boundary conditions. The corresponding adjoint operator consists of backward differences with Dirichlet boundary conditions. The resulting scheme reads:

$$u_x(i,j) = \begin{cases} u(i+1,j) - u(i,j) & \text{if } i < n_x \\ 0 & \text{if } i = n_x \end{cases}$$

$$u_y(i,j) = \begin{cases} u(i,j+1) - u(i,j) & \text{if } j < n_y \\ 0 & \text{if } j = n_y \end{cases}$$

$$\nabla \cdot \boldsymbol{y}(i,j) = \begin{cases} y_1(i,j) - y_1(i-1,j) & \text{if } i > 0 \\ y_1(i,j) & \text{if } i = 0 \\ -y_1(i-1,j) & \text{if } i = n_x \end{cases}$$

$$+ \begin{cases} y_2(i,j) - y_2(i,j-1) & \text{if } j > 0 \\ y_2(i,j) & \text{if } j = 0 \\ -y_2(i,j-1) & \text{if } j = n_y. \end{cases}$$

2.4.1 Warping Operator

The structure of the warping operator W^i depends on the interpolation scheme used to evaluate $\tilde{u}(i,j) = u(i + v_1(i,j), j + v_2(i,j))$. In our observation, a cubic interpolation scheme creates by far the best results.
The cubic interpolation scheme incorporates a 16-point neighborhood of the target position. This is given by all pairwise combinations $p_{kl} = (i_k, j_l)^T$ of

$$\begin{aligned} i_0 &= floor(i + v_1(i,j)) - 1, i_1 = i_0 + 1, i_2 = i_1 + 1, i_3 = i_2 + 1, \\ j_0 &= floor(j + v_2(i,j)) - 1, j_1 = i_0 + 1, j_2 = j_1 + 1, j_3 = j_2 + 1. \end{aligned}$$

Let us further denote

$$\begin{aligned} \mathcal{I}_{1D}(p_0, p_1, p_2, p_3, x) :=& (-\frac{1}{2}p_0 + \frac{3}{2}p_1 - \frac{3}{2}p_2 + \frac{1}{2}p_3)x^3 + (p_0 - \frac{5}{2}p_1 + 2p_2 - \frac{1}{2}p_3)x^2 \\ &+ (-\frac{1}{2}p_0 + \frac{1}{2}p_2)x + p_1 \end{aligned}$$

as the one-dimensional cubic interpolation at position x using values at position $p_0, \ldots, p_3$. Then, we obtain the bicubic interpolation of $\tilde{u}(i,j)$ as a combination of two one-dimensional problems:

$$\begin{aligned} \tilde{p}_0 &:= \mathcal{I}_{1D}(p_{00}, p_{01}, p_{02}, p_{03}, j + v_2(i,j)), \\ \tilde{p}_1 &:= \mathcal{I}_{1D}(p_{10}, p_{11}, p_{12}, p_{13}, j + v_2(i,j)), \\ \tilde{p}_2 &:= \mathcal{I}_{1D}(p_{20}, p_{21}, p_{22}, p_{23}, j + v_2(i,j)), \\ \tilde{p}_3 &:= \mathcal{I}_{1D}(p_{30}, p_{31}, p_{32}, p_{33}, j + v_2(i,j)), \\ u(i + v_1(i,j), j + v_2(i,j)) &\approx \mathcal{I}_{1D}(\tilde{p}_0, \tilde{p}_1, \tilde{p}_2, \tilde{p}_3, i + v_1(i,j)) \end{aligned}$$

Using this, one can build up the matrix W^i. For the row corresponding to the pair (i,j) the quantities $v_1(i,j)$ and $v_2(i,j)$ are static, thus we can calculate p_{kl}. Plugging this into the interpolation formula and regrouping it with respect to p_{kl} leads to an equation similar to $\sum_{k,l} \alpha_{kl} p_{kl}$. The weight α_{kl} is written at the column index corresponding to p_{kl}. In the boundary region we have to ensure that $i_0, j_0 \geq 0, i_3 \leq n_x$ and $j_3 \leq n_y$. If one of these conditions is violated, the row is set to zero. The adjoint operator to W^i is then simply the transposed matrix.

3 Parameters

In this section we give a short explanation of parameters and default choices for them, that are used in our experimental setup in Section 5. For param-

eters concerning the flow estimation we follow [20], where various factors in the optical flow calculation are analyzed and a robust baseline approach is deduced.
Stable stepsizes for the algorithm can be calculated using [15]. We recommend using diagonal matrices σ and τ instead of scalar values for better convergence speed. Thus, the respective parameter choices refer to each element.

- $nSales$ The number of scales in the coarse-to-fine pyramid is dynamically chosen, such that the image size on the coarsest scale is 10 pixels minimum in each direction.
- η The downsampling factor between subsequent levels is set to 0.8.
- $nWarps$ We do three warps per level of the pyramid.
- $sizeMed$ The median filtering is done in a 5x5 neighborhood.
- ϵ_v, ϵ_u The stopping tolerance for both, the u and the $\boldsymbol{v}$ problem is set to 10^{-6}.
- ϵ_{main} The stopping criterion for the main algorithm is set to 10^{-5}.
- $nRes$ The residual is only updated every 100 iterations, which is sufficient since the algorithm typically needs more that 1000 iterations and the calculation is computationally expensive.
- α Typically, we have for the regularization weight in the image sequence $\alpha \in [0.005, 0.1]$ depending on the noise in the input data. We usually do an initial calculation with $\alpha = 0.02$.
- β Since the flow part incorporates β and γ we propose to only adjust the ratio $\frac{\beta}{\gamma} \in [0.01, 0.1]$ and to start with a regularization weight for both flow components of $\frac{\beta}{\gamma} = 0.02$.
- γ The choice for γ regulates how strong the images are coupled by the flow component. In our experiments, $\gamma = 1$ is the usual choice.
- σ, τ We calculate for the $\boldsymbol{v}$-problem $\sigma_1^v = \sigma_2^v = \frac{1}{2}, \sigma_3^v = \frac{1}{\|\nabla \tilde{u}\|_1}, \tau_1^v = \frac{1}{4+|\tilde{u}_x|}$ and $\tau_2^v = \frac{1}{4+|\tilde{u}_y|}$. Analogously, we get for the u-problem $\sigma_1^u = \frac{1}{\|A\|_1}, \sigma_2^u = \frac{1}{2}, \sigma_3^u = \frac{1}{\|\mathcal{W}\|_1}$ and $\tau^u = \frac{1}{\|A^T\|_1+4+\|\mathcal{W}\|_1}$, where $\|\cdot\|_1$ for the matrix refers to the corresponding row-vector.

- $iterMainMax$ The maximum number of iterations for the main algorithm is set to 10 throughout our experiments.

4 Algorithm

The algorithm can be separated into three problems:

- A $\boldsymbol{v}$-part (1) that calculates the flow field between every two subsequent images.
- A u-part (2) that solves the reconstruction problem for the image sequence using the flow information.
- A main algorithm (3) that creates initial values and alternatingly calls the other two functions.

The main algorithm uses a stopping criterion based on the difference between consecutive values for u and $\boldsymbol{v}$ or alternatively stops after a defined number of iteration $iterMainMax$.
The flow part requires solving a series of minimization problems and, consequently, is the time consuming part of the main algorithm. The structure of the primal-dual algorithm allows parallelization of algorithm (1) and (2) on CPU or GPU to improve the runtime.

4.1 FlexBox

An easier approach to minimize the arising variational problems for u and $\boldsymbol{v}$ is given by **FlexBox** [7], a flexible primal-dual toolbox. The basic framework for this optimization toolbox is written in MATLAB, but the software comes with an optional C^{++} and CUDA module to massively enhance the runtime. **FlexBox** allows to formulate the primal variational problem term-by-term via MATLAB with only a few lines of code and calculates the solution in the fastest possible way. Creation of dual variables, calculation of step-sizes and transfer to the C^{++} or CUDA solver are fully automated. The full sourcecode can be downloaded from http://www.flexbox.im.

Algorithm 1 Subproblem for flow field $\boldsymbol{v}$

function PROBLEMV(u_1, u_2, α)
  Init $\boldsymbol{y}_1, \boldsymbol{y}_2, y_3, \boldsymbol{v}$ on coarsest scale
  for $s \leftarrow 1, nScales$ **do**
    Upsample $\boldsymbol{y}_1, \boldsymbol{y}_2, y_3, \boldsymbol{v}$ to current scale using bicubic interpolation
    for $w \leftarrow 1, nWarps$ **do**
      $\tilde{\boldsymbol{v}} \leftarrow \boldsymbol{v}$
      Calculate $\tilde{u}_2(x + \tilde{\boldsymbol{v}})$ and $\nabla\tilde{u}_2(x + \tilde{\boldsymbol{v}})$ w.r.t. s
      while $\epsilon_v < r_v^{pd}$ **do**
        $\boldsymbol{y}_1^{k+1} \leftarrow \frac{\boldsymbol{y}_1^k + \sigma\nabla\bar{v}_1^k}{\max(1, \frac{\|\boldsymbol{y}_1^k + \sigma\nabla\bar{v}_1^k\|_2}{\alpha})}$
        $\boldsymbol{y}_2^{k+1} \leftarrow \frac{\boldsymbol{y}_2^k + \sigma\nabla\bar{v}_2^k}{\max(1, \frac{\|\boldsymbol{y}_2^k + \sigma\nabla\bar{v}_2^k\|_2}{\alpha})}$
        $y_3^{k+1} \leftarrow \max(-1, \min(1, y_3^k + \sigma(\nabla\tilde{u} \cdot \bar{\boldsymbol{v}}^k + \tilde{u}_t)))$
        $v_1^{k+1} \leftarrow v_1^k - \tau(\nabla^T \cdot \boldsymbol{y}_1^{k+1} + \tilde{u}_x y_3^{k+1})$
        $v_2^{k+1} \leftarrow v_2^k - \tau(\nabla^T \cdot \boldsymbol{y}_2^{k+1} + \tilde{u}_y y_3^{k+1})$
        $\bar{\boldsymbol{v}}^{k+1} \leftarrow 2\boldsymbol{v}^{k+1} - \boldsymbol{v}^k$
        Update r_v^{pd} every $nRes$ iterations
      end while
      Apply median filter to $\boldsymbol{v}$
    end for
  end for
  return $\boldsymbol{v}$
end function

Algorithm 2 Subproblem for image sequence u

function PROBLEMU($f, \boldsymbol{v}, \alpha, \gamma$)
  Generate operators $\mathcal{A}, \bar{\nabla}, \mathcal{W}$
  while $\epsilon_u < r_u^{pd}$ **do**
    $y_1^{k+1} \leftarrow \frac{1}{\sigma+1}(y_1^k + \sigma(A\bar{u}^k - f))$
    $\boldsymbol{y}_2^{k+1} \leftarrow \frac{\boldsymbol{y}_2^k + \sigma\bar{\nabla}\bar{u}^k}{\max(1, \frac{\|\boldsymbol{y}_2^k + \sigma\bar{\nabla}\bar{u}^k\|_2}{\alpha})}$
    $y_3^{k+1} \leftarrow \max(-\gamma, \min(\gamma, y_3^k + \sigma\mathcal{W}\bar{u}^k))$
    $u^{k+1} \leftarrow u^k - \tau(A^T y_1^{k+1} + \bar{\nabla}^T \cdot \boldsymbol{y}_2^{k+1} + \mathcal{W}^T y_3^{k+1})$
    $\bar{u}^{k+1} \leftarrow 2u^{k+1} - u^k$
    Update r_u^{pd} every $nRes$ iterations
  end while
  return u
end function

Algorithm 3 Main iteration

function MAINFUNCTION(f, α, β, γ)
  $iteration \leftarrow 0$
  $\boldsymbol{v} \leftarrow 0$
  $u \leftarrow \text{problemU}(f, \boldsymbol{v}, \alpha, 0)$
  while $r_{main} > tol_{main}$ or $iteration < iterMainMax$ **do**
    $u_{old} \leftarrow u$
    $\boldsymbol{v}_{old} \leftarrow \boldsymbol{v}$
    for $i \leftarrow 1, nImages - 1$ **do**
      $v_i \leftarrow \text{problemV}(u_i, u_{i+1}, \beta, \gamma)$
    end for
    $u \leftarrow \text{problemU}(f, \boldsymbol{v}, \alpha, \gamma)$
    $r_{main} \leftarrow \|u - u_{old}\| + \|\boldsymbol{v} - \boldsymbol{v}_{old}\|$
    $iteration \leftarrow iteration + 1$
  end while
  return $(u, \boldsymbol{v})$
end function

4.2 Time-Continuous Model

We want to make clear that the given algorithm can be used to minimize the original time-continuous model from [6] with only a few modifications. By leaving out the coarse-to-fine pyramid and the intermediate warping steps in the flow problem $\boldsymbol{v}$ (i.e. setting $\tilde{\boldsymbol{v}} = 0$ in the evaluation of $\tilde{u}_2$ and $\nabla\tilde{u}_2$) we solve a classical optical flow problem. The sub-problem for the image sequence u can be adjusted by exchanging the warping operator $\mathcal{W}$ with a discretization of the equation $\nabla u_2 \cdot \boldsymbol{v} + u_2 - u_1 = K(u_1, u_2)^t$ with $K = [-Id, diag(v_1)\nabla_x + diag(v_2)\nabla_y + Id]$ does directly apply the calculated small-scale flows and has not to be changed at all.
The limitations from [6] of course apply and the model can only be used from sequences with motion of very small magnitude.

5 Experimental Results

We evaluate the algorithm on grayscale datasets from the Middlebury optical flow database [2], which contain at least five consecutive frames and have available ground truth flow fields. Each image is disturbed with Gaussian noise ($\mu = 0, \sigma^2 = 0.01$) to test both, the image denoising and the flow estimation capacities. We use static parameters presented in Section 3 throughout the evaluation.
To measure the image reconstruction quality, we calculate L^2-error, PSNR and SSIM (default MATLAB implementation) and average them with respect to the number of frames. To evaluate the calculated flow fields, we calculate endpoint error (EPE) [13] and angular error (AE) [8, 3] between our result and the given ground-truth flow field. In the Middlebury sequence, the ground truth field is only given for the flow between the fourth and fifth frame.
The resulting errors are listed in Table 1. Figure 1 gives a visual overview of the denoising capabilities of our algorithm. Both, calculated errors and visual impression indicate a robust and superior reconstruction capability for dynamic image sequences. Moreover, the zoom in Figure 1 shows nicely recovered edges, well reconstructed structures and less staircasing effects than comparable total variation regularization techniques. The error in the flow field is relatively high, compared to similar motion estimation techniques. However, one has to take into account that usual benchmarks are calculated

	Gove2	**Grove3**	**Hydrangea**	**RubberWhale**	**Urban2**	**Urban3**
SSIM	0.6920	0.6036	0.7412	0.8350	0.7103	0.6961
L^2-Error	3.70e-05	5.62e-05	3.85e-05	2.10e-05	2.91e-05	3.11e-05
PSNR	72.9669	69.2809	74.0326	79.2496	75.1465	74.4912
EPE	1.0290	1.8428	1.3262	0.6566	2.0765	2.2277
AE	0.1713	0.2115	0.1396	0.2514	0.1588	0.2395

Table 1: Image reconstruction and flow estimation errors for sequences (manually disturbed with Gaussian noise) from the Middlebury database [2].

on clean input images.

6 Conclusions

In this work, we propose a detailed implementation of the work presented by Burger, Dirks and Schönlieb [4] and extend their time-continuous model for small-scale flows to arbitrary large-scale motion between consecutive frames. The arising variational problem is minimized using a primal-dual method and the operator-discretization is explained in detail. We present concrete parameter choices and give a pseudo-code of the implemented algorithm. Finally, a short numerical evaluation based on data from the Middlebury optical flow database is presented.
An obvious possible extension would be to deal with color-images. The color-channels can be coupled here to improve the results. Moreover, an implementation to higher spatial dimensions is possible, which causes an increased runtime, and also may require a different motion model (e.g. mass-preservation) on the other hand.

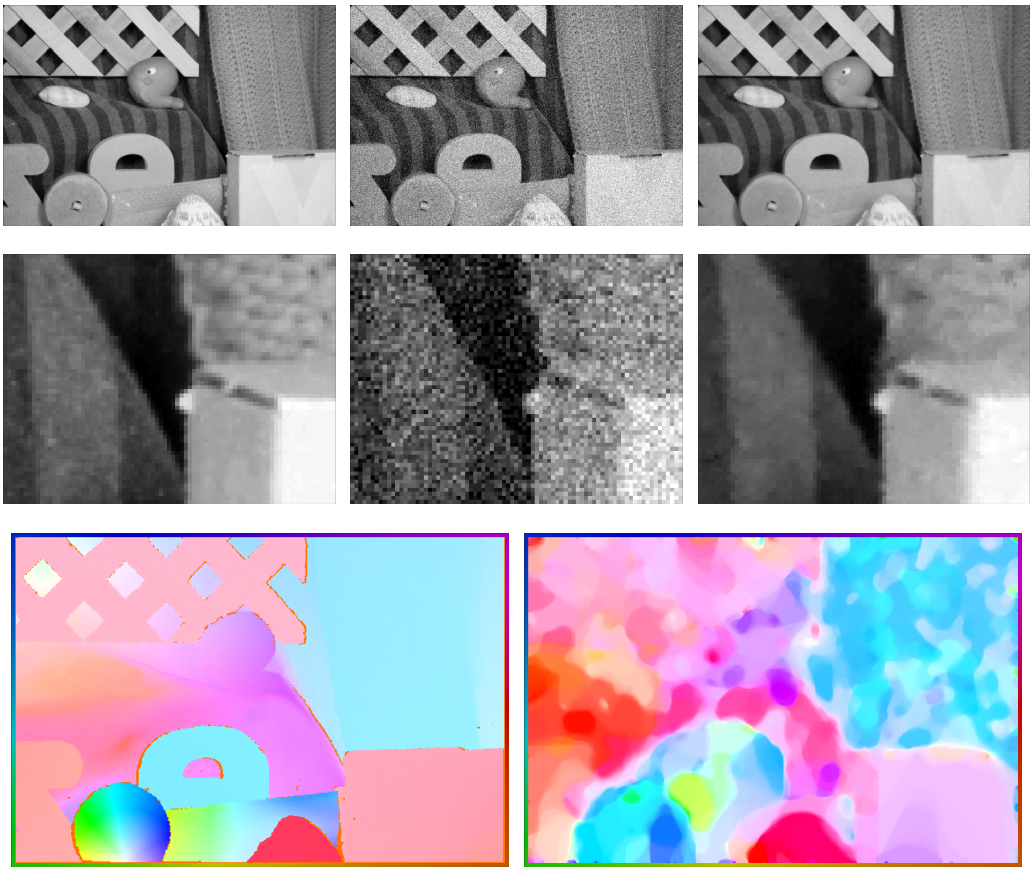

Figure 1: Frame 4 of the Rubber Whale sequence from [2]. The top row shows the full input image, disturbed counterpart and reconstruction. The middle row contains a zoom into the bottom right part of the image. In the bottom row we see ground truth and estimated flow field.

References

[1] G. Aubert, R. Deriche, and P. Kornprobst, *Computing optical flow via variational techniques*, SIAM Journal on Applied Mathematics, 60 (1999), pp. 156–182.

[2] S. Baker, D. Scharstein, J. Lewis, S. Roth, M. J. Black, and R. Szeliski, *A database and evaluation methodology for optical flow*, International Journal of Computer Vision, 92 (2011), pp. 1–31.

[3] J. L. Barron, D. J. Fleet, and S. S. Beauchemin, *Performance of optical flow techniques*, International journal of computer vision, 12 (1994), pp. 43–77.

[4] M. Burger, H. Dirks, and C.-B. Schönlieb, *A variational model for joint motion estimation and image reconstruction*, arXiv preprint arXiv:1607.03255, (2016).

[5] A. Chambolle and T. Pock, *A first-order primal-dual algorithm for convex problems with applications to imaging*, Journal of Mathematical Imaging and Vision, 40 (2011), pp. 120–145.

[6] H. Dirks, *Variational Methods for Joint Motion Estimation and Image Reconstruction*, PhD thesis, WWU Münster, 2015.

[7] ——, *A flexible primal-dual toolbox*, arXiv preprint, (2016). http://www.flexbox.im.

[8] D. J. Fleet and A. D. Jepson, *Computation of component image velocity from local phase information*, International Journal of Computer Vision, 5 (1990), pp. 77–104.

[9] T. Goldstein, E. Esser, and R. Baraniuk, *Adaptive primal-dual hybrid gradient methods for saddle-point problems*, arXiv preprint arXiv:1305.0546, (2013).

[10] B. K. Horn and B. G. Schunck, *Determining optical flow*, in 1981 Technical Symposium East, International Society for Optics and Photonics, 1981, pp. 319–331.

[11] T. Kösters, K. P. Schäfers, and F. Wübbeling, *Emrecon: An expectation maximization based image reconstruction framework for emission tomography data*, in Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2011 IEEE, IEEE, 2011, pp. 4365–4368.

[12] D. Mitzel, T. Pock, T. Schoenemann, and D. Cremers, *Video super resolution using duality based TV-L1 optical flow*, in Pattern Recognition, Springer, 2009, pp. 432–441.

[13] M. Otte and H.-H. Nagel, *Optical flow estimation: advances and comparisons*, in Computer VisionECCV'94, Springer, 1994, pp. 49–60.

[14] N. Papenberg, A. Bruhn, T. Brox, S. Didas, and J. Weickert, *Highly accurate optic flow computation with theoretically justified warping*, International Journal of Computer Vision, 67 (2006), pp. 141–158.

[15] T. Pock and A. Chambolle, *Diagonal preconditioning for first order primal-dual algorithms in convex optimization*, in 2011 International Conference on Computer Vision, IEEE, 2011, pp. 1762–1769.

[16] T. Pock, D. Cremers, H. Bischof, and A. Chambolle, *An algorithm for minimizing the Mumford-Shah functional*, in Computer Vision, 2009 IEEE 12th International Conference on, IEEE, 2009, pp. 1133–1140.

[17] L. I. Rudin, S. Osher, and E. Fatemi, *Nonlinear total variation based noise removal algorithms*, Physica D: Nonlinear Phenomena, 60 (1992), pp. 259–268.

[18] A. Sawatzky, C. Brune, F. Wubbeling, T. Kosters, K. Schafers, and M. Burger, *Accurate EM-TV algorithm in pet with low SNR*, in 2008 IEEE Nuclear Science Symposium Conference Record, IEEE, 2008, pp. 5133–5137.

[19] J. Shen and T. F. Chan, *Mathematical models for local nontexture inpaintings*, SIAM Journal on Applied Mathematics, 62 (2002), pp. 1019–1043.

[20] D. Sun, S. Roth, and M. J. Black, *A quantitative analysis of current practices in optical flow estimation and the principles behind them*, International Journal of Computer Vision, 106 (2014), pp. 115–137.

[21] M. Unger, T. Pock, M. Werlberger, and H. Bischof, *A convex approach for variational super-resolution*, in Pattern Recognition, Springer, 2010, pp. 313–322.

[22] Y. Wang, W. Yin, and Y. Zhang, *A fast algorithm for image deblurring with total variation regularization*, 2007.

[23] A. Wedel, T. Pock, C. Zach, H. Bischof, and D. Cremers, *An improved algorithm for TV-L1 optical flow*, in Statistical and Geometrical Approaches to Visual Motion Analysis, Springer, 2009, pp. 23–45.

[24] C. Zach, T. Pock, and H. Bischof, *A duality based approach for realtime TV-L1 optical flow*, in Pattern Recognition, Springer, 2007, pp. 214–223.